# Nebula Testing Initiative

## An FY11/12 Collaboration: OCIO & SMD

March 28, 2012


**Mike Little**
**m.m.little@nasa.gov (202)358-0954**

# Presentation Outline

- Background
- Nebula Testing
- Conclusions
- The Way Forward

# Background: Key HEC Questions

- What are the comparative advantages of new technology over legacy?
  - How does this technology perform its function compared to the current technology?
    - Comparisons by literature
    - Benchmarks
    - Experiments
  - How hard is it for the Science & Engineering Community to adopt it
- What business model works best?

# Background:
# S&E and Business Computing Differ

- Commercial Off-the-Shelf Software
  - MATLAB, IDL, plotting tools, ProEngineer, PATRAN, etc.
- Custom Science Code
  - Developed by Science Community
    - Is not likely to have a strong SW Engineering basis
    - Dependent upon public domain or open source libraries
    - Compiler – Commercial vs. Open Source
  - Processor Performance (hardware and operating system)
    - Instruction Throughput
    - Instruction faults and errors
    - When running, occupies 80% processor cycles
  - Communications
    - Interprocessor communications
      - Within Node
      - External to Node, but local within system
      - External to
    - File Access
      - Scratch space on local disks
      - High throughput, low latency to keep processors executing
  - Data Volumes (primarily demands fast communications)
    - Large input (no need to preserve)
      - Staged locally vs. Accessed external to system
    - Large output to local disk, but not externally
    - Large output to be staged onto external storage for later evaluation and potential retention
  - Accessibility and Credentials
    - External users
    - Collaboration includes non-US users located outside NASA domain

***HHEC/OCIO is performing an S&E computing needs assessment to improve understanding of mission directorate computing needs, processes and workflows. From this, we can tune the computing portfolio to be most responsive.***

# Nebula Testing Project Initiation

- Culmination of 2 years of Development by Ames CIO and NASA CIO
  - OpenStack as an open source Cloud package
    - NASA participation motivated others to join
  - Initially OCIO saw target customer HECC
- SMD Leased Nebula for 5 months to evaluate its value in Science & Engineering Computing
  - Testing opportunity started on 7/5

# Nebula Testing Overall Approach

- Phase 1A Characterize Nebula through Testing
  - Debug the system using science codes
  - Evaluate Performance roadblocks
  - Gain Experience with Nebula to help Projects
  - Quantify performance compared to Amazon, HEC
- Phase 1B – Evaluate Nebula's Utility to Projects through Testing
  - Qualitative, anecdotal experience by Projects
  - Quantitative Performance measured
    - Compare to the normal environment
- Phase 1 Milestone: Review Board Evaluation
- Phase 2 Planned to use Nebula as an HECC asset

# Nebula Testing Results

- Phase 1 Completed but NOT as we Expected
  - Nebula characterization: Ethernet Performance is crippling
  - Amazon Cloud (HPC and GovCloud) characterization: Inter-node Comm is marginally acceptable
  - Microsoft Azure Cloud Characterization: Windows Only and Comm slow
  - Phase 1A required correcting Nebula deficiencies before users introduced onto system
  - Phase 1B revealed many limitations on User acceptability of Nebula
  - User Acceptance of Amazon Cloud far better
- Phase 1 Decision Point was Cancelled
  - Evidence overwhelming that Nebula was insufficiently mature to supply Production-quality service
- Phase 2 in Two Thrusts – Pay as You Go
  - Production Computing using Commercial Sources
  - NASA Testbed to further mature technologies and give users exposure

# Conclusions from Nebula Testing

Overall

- The state of Cloud Computing is rapidly changing
- We need a set of benchmarks to compare Cloud Performance
- Platform-as-a-Service is more Useful to S&E than Infrastructure-as-a-Service
- The Business Interface to Commercial Clouds is VERY Important
- Commercial Services offer better Cloud elasticity and economics and improvement
- Cloud Security Certification is VERY Important
- Nebula does not currently offer operational capabilities over AWS
- Nebula implementation needs further maturation to become operational
- Migration to Cloud Computing is Accelerated by availability of experienced User Support

# The Way Forward
# S&E Applications on Clouds
## What Works and What Doesn't

- Characteristics of Applications Well Suited to HPC Cloud
  - New Code is easier to instantiate than re-hosting code designed for HEC
  - All work performed on a single node, generally by 4 cores
    - Embarrassingly parallel computations, for example work well on many nodes
  - Network bandwidth < 10Mbps needs for input/output data
  - Compile with Intel commercial compilers
  - Memory usage under 2 GB per core
  - Application is bursty or at infrequent intervals
    - Remain logged off otherwise
  - VM instance is highly configured by end-user, differently than existing HEC systems
  - Applications requiring collaboration with non-US citizens
- Characteristics of Unsuited Applications
  - High data flow off-cloud
  - Legacy code which has not been re-hosted recently (not very portable to anything else)
  - Legacy code which uses out of date or specialized compilers (best is Intel compiler)
  - Application needs to run continuously (not economic)
  - Heavy dependence upon MPI, large shared memory space, or multicast networking
- User-team Requirements
  - Strong SysAdmin skills available
  - Strong knowledge/experience in Operating Systems, compilers and libraries
  - Can live within boundaries of customizable commodity environments
  - Funded at least at a modest level
  - Understanding of the requirements for verifying scientifically identical output products

# The Way Forward
# Mission S&E Computing in the Cloud

- Each Center CIO CTO is conducting their own effort
  - Pay as you go
  - JPL Business Interface not available in non-CalTech NASA
- HECC Project at ARC is implementing a supplement to Pleiades
  - Initially available at HECC discretion to offload surge volume on queues
  - Later, will allow users to buy time to accelerate job execution
- User-team Requirements
  - Strong SysAdmin skills available
  - Strong knowledge/experience in Operating Systems, compilers and libraries
  - Can live within boundaries of customizable commodity environments
  - Funded at least at a modest level
  - Understanding of the requirements for verifying scientifically identical output products

# The Way Forward
# JPL Use of Cloud Computing

- The CIO's CTO, Tom Soderstrom, has teamed with several of the mission organizations to experiment with Commercial Cloud Computing. They have created a range of options for using Cloud Computing, including Google, Amazon and a private arrangement with Lockheed Martin. The private cloud supports ITAR and Sensitive But Unclassified data.
- The JPL CIO's office has crafted a suite of procurements to order the services and a central billing function so that usage is paid appropriately. These procurements also provide a means of retrieving the data stored in the Cloud environment, should re-hosting be desired. They involved auditors in their design to ensure that the system would support external review, forensics and auditing of security, information security, and finance.
- They have also created a custom tool, CHASM, to help customers define their needs and to select the proper service offering, which is then manually provisioned by System Administrators and delivered to the customers within a very short time. It may need expansion as finer-grained distinctions between classes of computing. An ORR for the JPL capability is scheduled for next week.
- One application demonstrated was the ongoing mission data processing for the MER Program. The original equipment used for this effort when the Rovers first landed have become obsolete and difficult to maintain. Teaming with the CIO's office, the Project evaluated and then migrated their processing to this Cloud Computing environment and has been able to eliminate a significant re-investment in hardware

# The Way Forward
# A NASA Cloud Testbed

- Purpose: Support Mission Directorates
  - Evaluate improvements to cloud software stacks for S&E applications
  - Provide assistance to NASA cloud software developers
  - Assist appropriate NASA S&E users in migrating to clouds
- Operate cloud instances at ARC and GSFC
- Limited Funding under HEC Program as a Technology Development
- Liaison with other NASA centers and external S&E cloud users
- Partnership between OCIO and HECC
  - OCIO contributes:
    - Existing hardware
    - Liaison with and expertise from OpenStack
  - SMD contributes:
    - Production computing expertise
    - Supplemental hardware
    - O&M and user support
    - Expanded documentation and configuration management

# Many Thanks to Participants

## Karen Petraska of OCIO

**Test Element**

- Rupak Biswas
- Jerry C. Yan
- Piyush Mehrotra
- Bryan Biegel
- Dan Duffy
- Hoot Thompson
- Tom Soderstrom
- Khawaja Shams
- Phil Milstead
- Joe Bredekamp
- Al Settell

**Nebula Element**

- Tsengdar Lee
- James Williams
- Keith Shackleford
- Demetra Souliotis
- Chris Kemp
- Adrian Gardner
- Lon Gowan

**Project Participants**

- Ashutosh Limaye
- Dana Hammond
- Karl Loutinsky
- Ed Masuoka
- Rama Nemani
- Steve Maher
- John Schnase
- Max Lui
- Andrew Molthan
- Michelle Eshow
- Doris Pan
- Hamid Oloso
- Anning Cheng
- Kuan-man Xu
- Sharon Rodier
- Brian Getzewich
- Amanda O'Connor

# Backup Materials

# Background: HEC Objectives

- Do more science within the same budget
  - New technologies MUST be usable and acceptable to science communities
- Provide Supercomputing Computational Capacity and Capability to appropriate Agency Projects
- Improve effective use of supercomputing by the Science and Engineering users
  - Overall Architecture
  - Processors
  - Data delivery to the code
  - Human factors supporting Science & Engineering Computation
    - Integrated development environments
    - Support to Code calibration, validation
    - Community ability to use the results of HEC efforts

# Background: HEC Studies System Architectures

- Use of emerging commercial technology
  - Cloud Computing Capabilities
  - Cloud Storage
  - Hybrid cluster-GPU systems
- Use of unique or specialized technologies
  - Hybrid cluster-FPGA

# Nebula Testing Overall Approach

- Phase 1A Characterize Nebula through Testing
  - Debug the system using science codes
  - Evaluate Performance roadblocks
  - Gain Experience with Nebula to help Projects
  - Quantify performance compared to Amazon, HEC
- Phase 1B – Evaluate Nebula's Utility to Projects through Testing
  - Qualitative, anecdotal experience by Projects
  - Quantitative Performance measured
    - Compare to the normal environment
- Phase 1 Milestone: Review Board Evaluation
  - Cancelled due to overwhelming evidence of system inadequacy
  - Phase 2 completely re-vamped based on experience
- Phase 2 Plan
  - Meet Operational needs from commercial sources
  - Partner with OCIO on an Agency Testbed

Red Indicates Revised Plan

# Nebula Testing Schedule

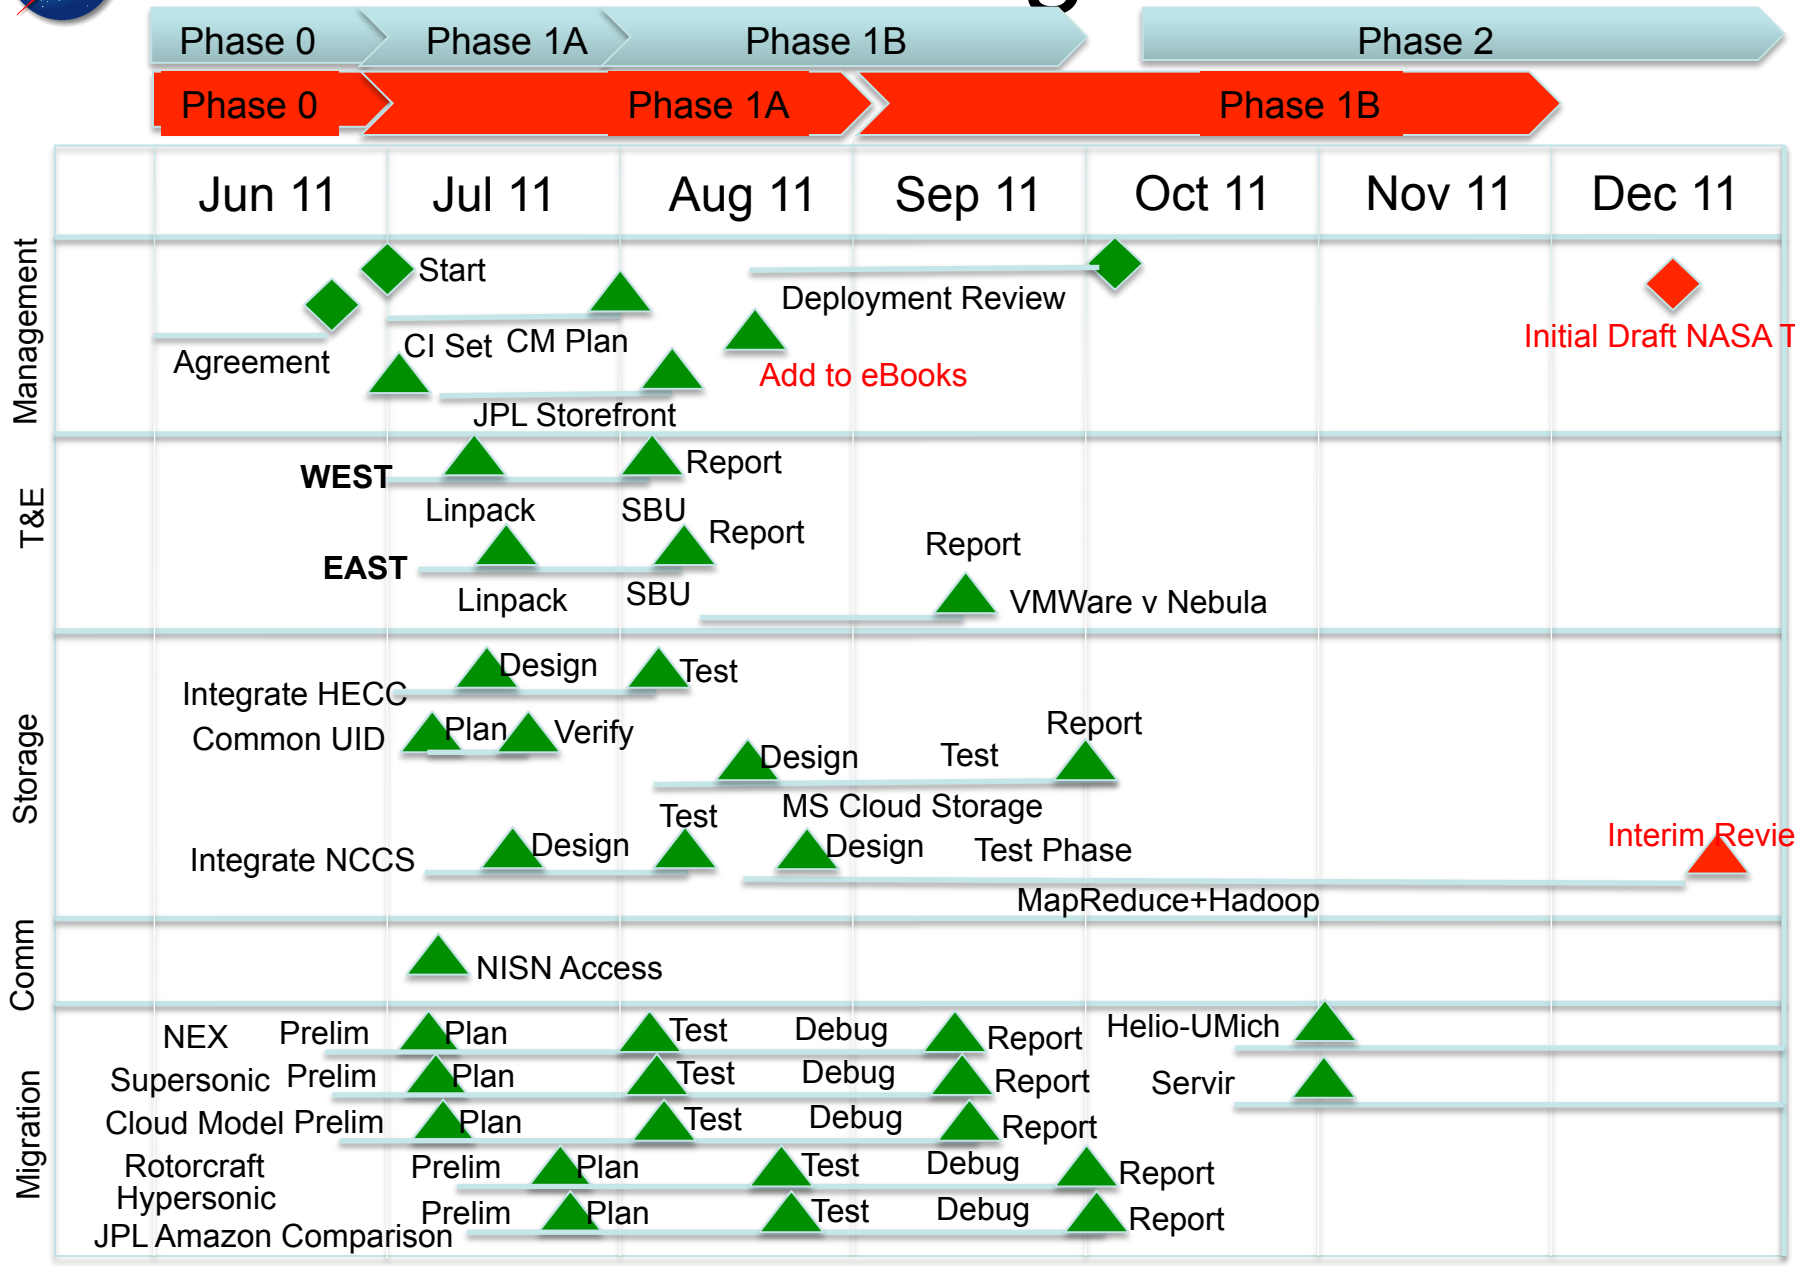

# Nebula Testing Results

- Phase 1 Completed but NOT as we Expected
  - Nebula characterization
    - Processor performance as good as Pleiades
    - I/O as good on the Node as desktop computers if data is on same node
    - Internode comm significantly poorer than Amazon
    - Considerable improvement in I/O and Comms from break-fix testing of Benchmark codes
  - Amazon Cloud (HPC and GovCloud) characterization
    - Processor performance as good as Pleiades
    - Internode comm reduces effective performance of related applications
      - Compared to 10GigE Ethernet, not fiberchannel or infiniband
  - Microsoft Azure Cloud Characterization
    - Windows Only at this stage
    - Similar comm problems
  - Phase 1A required correcting Nebula deficiencies before users introduced onto system
    - Performance was unacceptable for most S&E Applications
    - Many problems with User Acceptability
  - Phase 1B revealed many limitations on User acceptability of Nebula
  - User Acceptance of AWS far better
- Phase 1 Decision Point was Cancelled
  - Evidence was overwhelming that Nebula was insufficiently mature to supply Production-quality service
- Phase 2 in Two Thrust – Pay as You Go
  - Production Computing using Commercial Sources
  - NASA Testbed to further mature technologies and give users exposure

# Conclusions from Nebula Testing

## Overall

- Collectively, we learned a lot
  - Better definition of requirements for S&E computing in the Cloud Computing Environment
  - Classes of applications that could potentially benefit from cloud services, as they are today
  - Performance/supportability characteristics of Cloud Computing, in general, and Nebula, in particular
  - Testers have tuned up their usual tools to evaluate alternative Cloud Computing solutions
- Nebula does not currently offer operational capabilities over AWS
  - AWS performance, stability and reliability are better than Nebula
  - NASA HEC systems perform better than Nebula or AWS for NASA HPC applications
    - Lrgely because of the maturity of the tuning and the specialized hardware
  - Enormous public clouds, like AWS, with many customers have benefits over small private clouds
    - Smaller private clouds do not demonstrate the unit cost of large, multi-customer private clouds
    - Small, isolated clouds do not offer comparable quick response to large growth in demand
    - Large, multi-customer clouds can sustain re-investment in development to improve performance and capability
- Platform-as-a-Service is more Useful to S&E than Infrastructure-as-a-Service
- Nebula implementation needs further maturation to become operational
  - OpenStack implementation needs additional development in order to provide acceptable level of performance to more than a small number of S&E applications
  - Enhanced I/O infrastructure is needed for efficient and consistent storage and access to both persistent and transient data
  - Support model and SLA needs transformation to provide adequate and efficient support to Users
    - Administration as well as monitoring and communicating to users as customers
  - Mechanism for provisioning and tracking pay as you go service is needed
- Re-examine Nebula West hardware for supportability and unpredictable performance
- Migration to Cloud Computing is Accelerated by availability of experienced User Support

# Conclusions from Nebula Testing:

## NASA can leverage Cloud Computing in S&E Computing

- A NASA private cloud could offer advantages over public cloud (i.e.,AWS)
  - Control and changing (on demand) the configuration of the cloud software and hardware
  - Support experiments with hardware not available through commercial cloud providers
    - Infiniband, dynamic bare metal testing, non-standard virtual machine environments
  - Provide easy end-user access to cloud services with a Moderate Security Rating
    - In AWS, end users must implement some of the security protective measures and obtain certification
    - An Agency-wide value-added service supplementing AWS could do the same thing
    - Permit users to perform SBU computations without significant delays in obtaining approvals
  - Testbed to further expand capabilities of open source cloud software that industry would not
- Fully develop the capabilities of Cloud Computing to support S&E computation
  - Bursting from a private cloud onto multiple, competing clouds, both public and private
    - Enable a small private cloud to act as an entry point for accessing (appropriately) public cloud capacity
    - Encourage ongoing competition among cloud service providers
  - Low latency, high bandwidth access to remote storage without large in-cloud staging
  - Improve linkage between cloud services and existing storage environments at DAACs, modeling centers, lab data archives to avoid duplicating or delays due to staging data local to processing
  - Substantially more efficient ethernet drivers to maximize I/O performance
  - Internal File storage performance issues due to NFS implementation
- NASA needs continued availability of a development environment
  - Interactions between cloud system dictates two separately managed clouds
  - Provide a development and test environment for developers under AIST and other Agency efforts
  - Benchmark improvements to evaluate public and private cloud offerings in S&E context

# Conclusions from Nebula Testing:

## Further development of OpenStack software is beneficial to the advancement of cloud computing capabilities

- Multiple contributors maintain fast-paced development of new features and capabilities
  - Improve rate at which functional improvements are more universally available
- Maintain competitive array of cloud providers
  - Needs an array of software stacks (e.g., OpenStack, OpenNebula)
  - Similar to OS-wars (Mac, linux, unix vs. DOS/Win)
- Reduce the barrier to entry by small businesses to supply Open Source service
- Nebula evolution could show the path for improvements to other Cloud Systems
  - Demonstrate feasibility
  - Demonstrate the value to end-users in competitive markets

# Considerations in SMD for Funding FY12

- SMD funding of Nebula would be to satisfy Mission Directorate computational needs
  - Dependent upon actual, measured capability and its utility
  - An operational computer system
  - Funding for general purpose admin or training separate from this
- Dependent upon Nebula's ability to support doing more Science and Engineering within Budget
  - Clean-up technical problems discovered during testing
  - Business Model, including a viable, easy to use Storefront
- User acceptability is important
  - May need to fund willing users to convert code to migrate into cloud
  - Moderate Rating for Cybersecurity
  - File I/O and staging

# Science Community

Do More Science within Budget

- Research Projects lead by PI's
  - PI's competitive for awards
    - Responsible for meeting their commitments and publishing papers
  - Often Geographically dispersed teams
  - Under a lot of pressure to produce results
  - Trade-offs among contractors, post-docs, IT hardware, collaboration with other groups
  - Time to discovery and their own labor are the drivers
  - Collaboration with non-NASA scientists,often non-US citizens, is essential
    - Universities
    - Collaboration with scientists abroad
- HEC/OCIO Needs Assessment in Progress to improve definition of Science Computing Needs
  - How to do more science within budget

# Engineering Community

Project Completion

- Focus is on completing a system project
  - Public-private partnerships growing
  - Space-Act Agreements with companies
  - Commercial launch vehicles
  - Hosted payload missions
- Computing needs
  - Commercial engineering tools running on Dept or Project servers
  - Many research engineering codes behaving similar to Custom Science Codes

# Phase 1 Testing Strategy

## Where are the edges of utility?

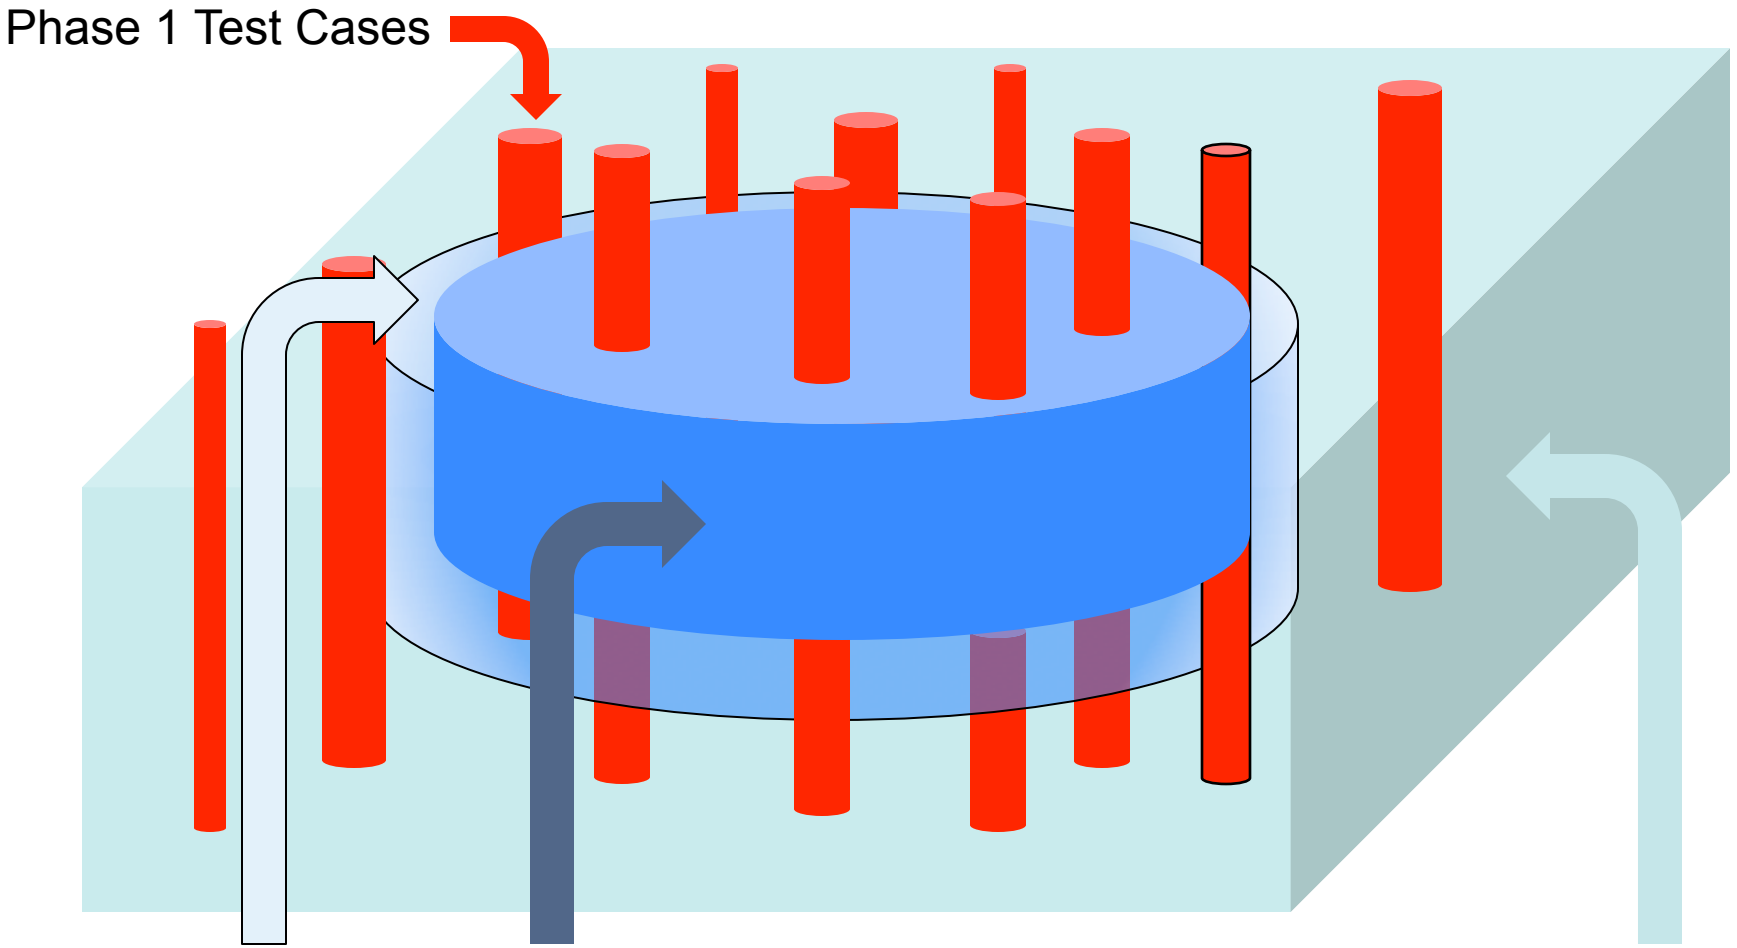

# Characterization Testing

- Linpack Benchmark
  - A dense set of linear equations
    - 64 bit floating point operations
    - operations are either addition or multiplication
    - Embarrassingly parallel
  - Ancillary Effects
    - Loads processors similarly to custom science codes with high instruction throughput
- HEC Standard Billing Unit Benchmark
  - Used to allocate, fairly, supercomputing assets
  - XX different codes, taken as a whole, represent full range of work done on HEC assets

# Phase 1A: Linpack Test Elements

A Dense Set of Linear Equations solved in FORTRAN

| Test Element | Purpose in Nebula | Key Characteristics | Conclusion |
|---|---|---|---|
| N=100 | Ability to run S&E Codes, fully loading processor, HiVAC | No changes to source code, 100x100 matrix, compiler optimization permitted | HIVAC and processing OK, off-node comm initially inadequate but much better |
| N=1000 | Ability to run S&E Codes, fully loading processor, HiVAC | Scaling | Too few Nebula processors available for comparison to capacity systems |
| Highly Parallel Computing Benchmark | Compare to other known systems | Highly parallel, uncoupled | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

# Phase 1A: SBU Test Elements

| Test Element | Non-Linearity | Inter-Grid Coupling | Off-Node Comm |
|---|---|---|---|
| **Enzo**: An adaptive mesh refinement, grid-based hybrid code designed to simulate cosmological structure formation (University of California, San Diego, et al.). | | | File I/O output<br>Latency:<br>Volume: |
| **FUN3D**: An unstructured computational fluid dynamics (CFD) code for analysis, adjoint-based error estimation, mesh adaptation, and aerospace design optimization extending into the hypersonic regime (NASA Langley Research Center). | MPI | MPI | File I/O output<br>Latency:<br>Volume: |
| **Goddard Earth Observing System Model, Version 5 (GEOS-5)**: The atmospheric general circulation model from the GEOS-5 system of models that also integrates an atmospheric analysis to support climate and weather prediction, data analysis, observing system modeling and design, and basic research (NASA Goddard Space Flight Center, et al.). | | Close coupling among processors | File I/O<br>Latency:<br>Volume: |
| **OVERFLOW**: A CFD program for solving complex flow problems that is widely used for designing launch and re-entry vehicles, rotorcraft, ships, and commercial aircraft (NASA Langley Research Center). | | | File I/O<br>Latency:<br>Volume: |
| **USM3D**: An unstructured mesh code for calculating flows over complex geometries that is often used to analyze aerodynamic flow of aerospace vehicle designs (NASA Langley Research Center). | | | File I/O<br>Latency:<br>Volume: |
| **Weather Research and Forecasting (WRF) Model**: A next-generation mesoscale numerical weather prediction system designed to serve both operational forecasting and atmospheric research needs (National Center for Atmospheric Research, et al.). | | | File I/O<br>Latency:<br>Volume: |

# Phase 1B Sample Trial Projects

| Project | Contact | Center | Mission | Description |
|---|---|---|---|---|
| SERVIR | Ashutosh Limaye | MSFC | SMD Applied Sciences | Data-driven Weather Forecasting for decision makers |
| MODAPS | Ed Matusoka | GSFC | SMD EOS-DIS | Remote sensing data products |
| 3D Radiation | Brian Getzewich | | SMD | Earth's Radiation Budget |
| NEX | Rama Nemani | ARC | SMD (ESD) | Model Access for non-model experts |
| WISE | Stephen Maher | GSFC | SMD Astrophysics | Wide-field Infrared mapping of entire sky. Data reduction. |
| ATC | | | | |
| Cloud Modeling | Kuan-man Xu Anning Cheng | LaRC | SMD (ESD) | Multi-scale modeling framework to represent cloud physical processes |
| AIRS | Long Pham | GSFC | SMD (ESD) | Remote Sensing data products |
| MATLAB | Sharon Rodier | LaRC | SMD (ESD) | Data Analysis using MATLAB hosted in both Nebula and Amazon |
| IDL in the cloud | Brian Getzewich | LaRC | SMD (ESD) | Data analysis using IDL hosted in both Nebula and Amazon (Commercial) |
| GEOS Cube Sphere | Dan Duffy | GSFC | SMD (ESD) | Data assimilation |
| WRF | Doris Pan | GSFC | SMD (ESD) | Commonly-used weather model |
| GISS Model E | Hamid Oloso | GSFC | SMD (ESD) | Climate model tuned to run on multi-core nodes with minimal inter-node communications |

# Nebula Testing Schedule

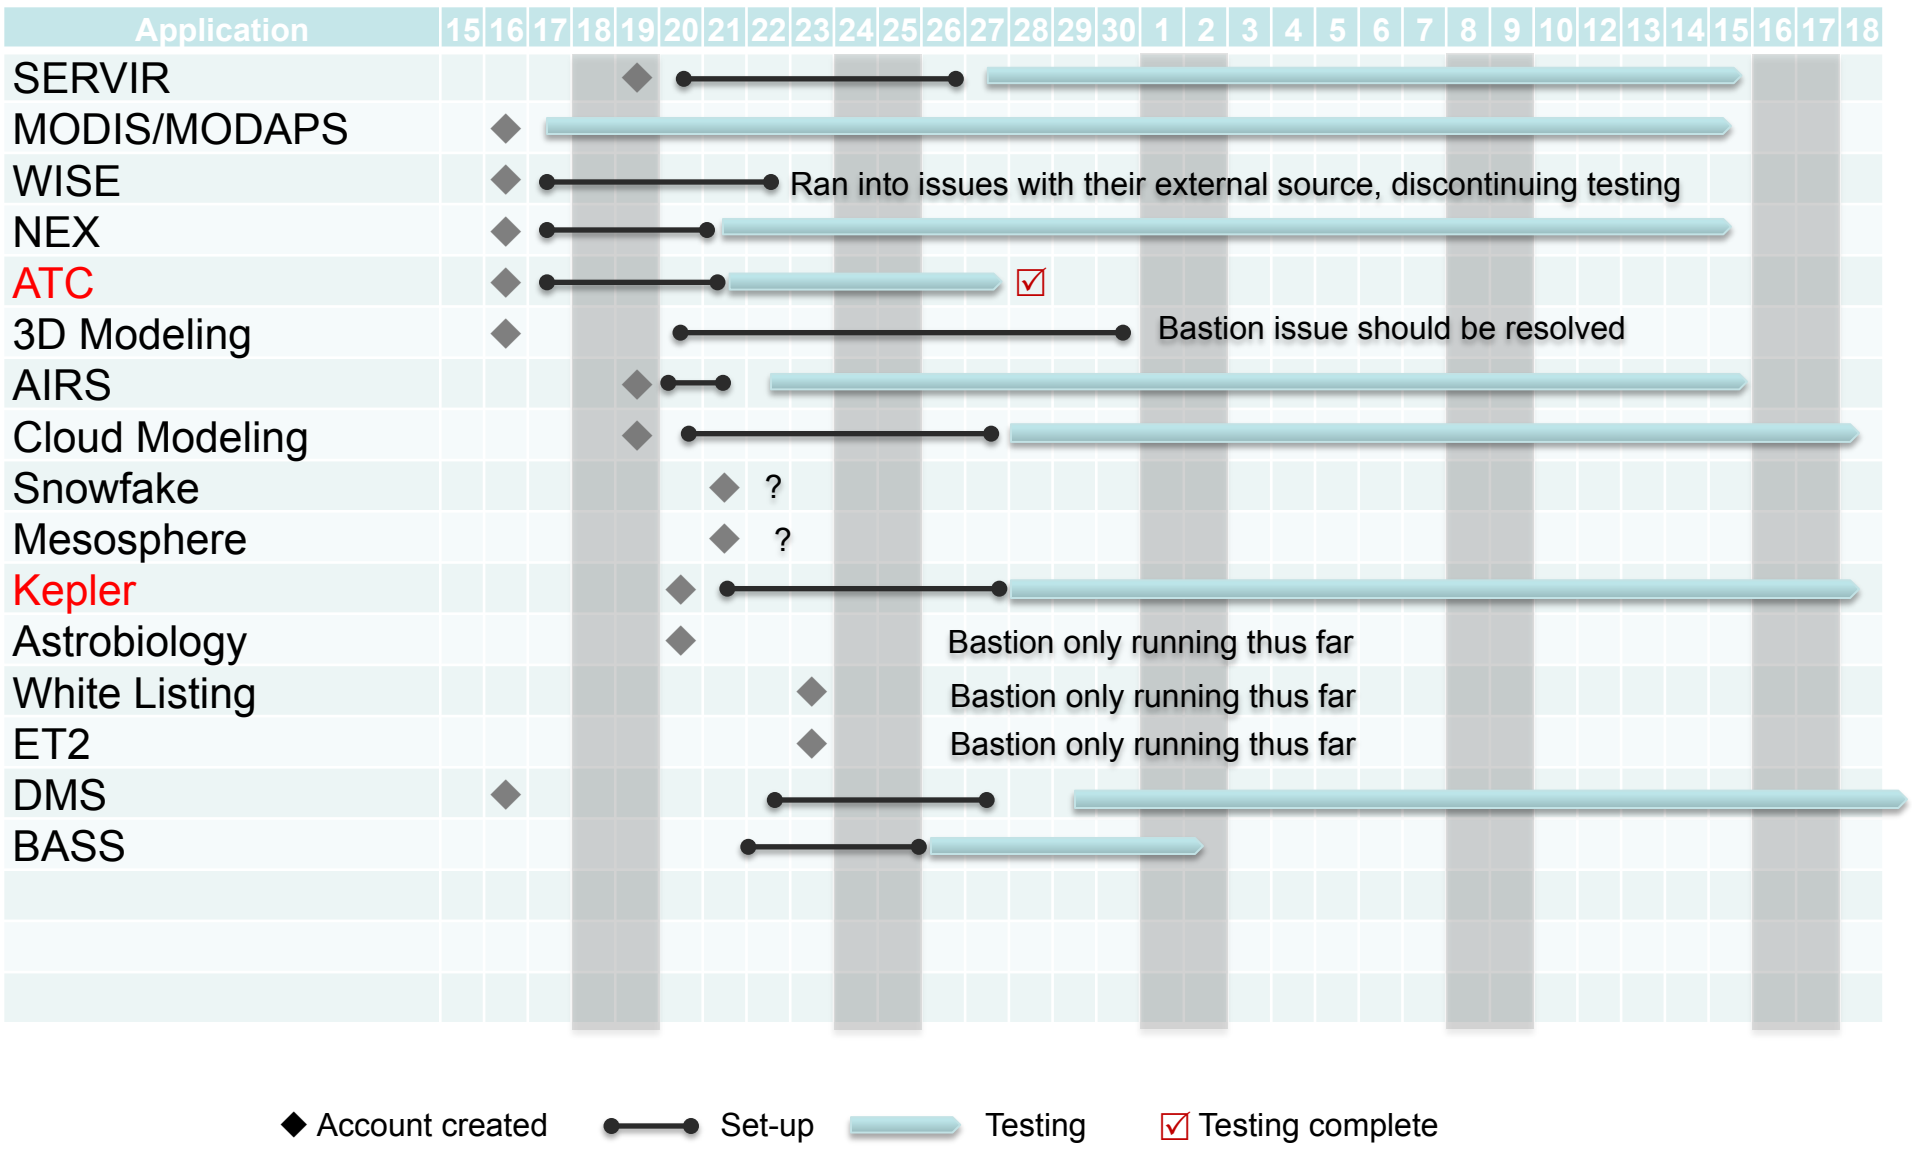

# Classes of Nebula Use Cases

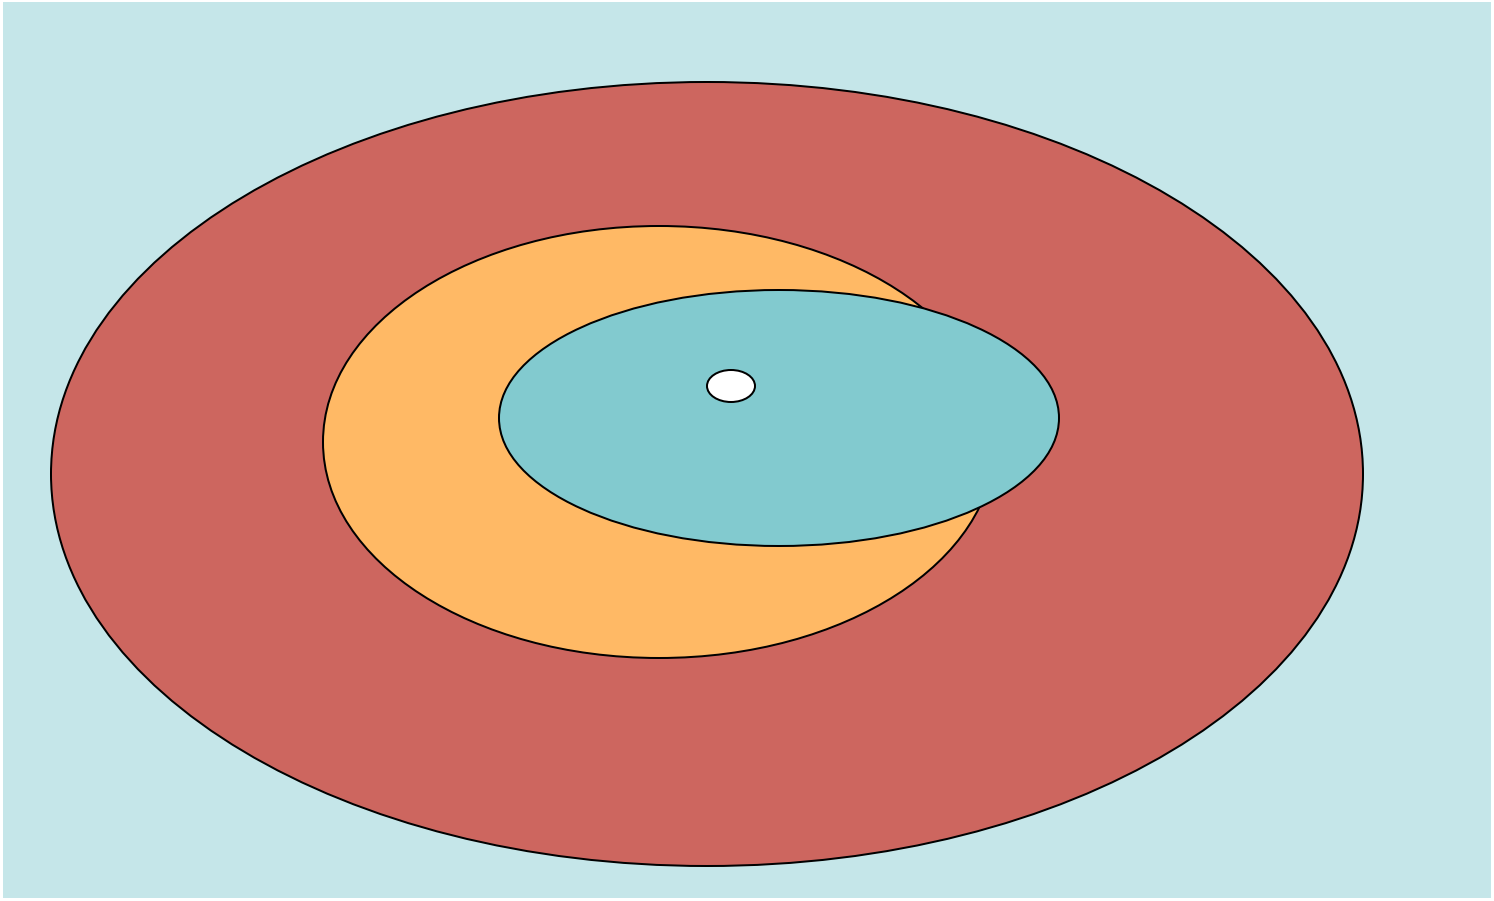

# History of NASA improving IT industry

- 1960's IBM s/360 Operating System
- 1980's Cray Unicos to UNIX
- 1990's SGI
- Currently
  - Cray uRiKA
  - OpenStack

# HECC Proposal for Cloud Testbed

*Support the enhancement of Cloud Computing capabilities for NASA S&E tasks*

- **Approach:**
  - Install, operate, and manage in-house cloud test-bed
  - Identify S&E benchmarks for evaluating cloud software stacks
  - Install at least two cloud software stacks, e.g., OpenStack, Open Nebula, Eucalyptus, etc.
  - Conduct performance evaluation using the benchmarks and identify specific and urgent performance issues
  - Provide feedback to cloud software developers
  - Coordinate efforts with interested parties at other Federal agencies, universities and commercial outfits including cloud developers, providers, and assessment teams